

Grant Number: NAG-1-2168

**An Advanced Hierarchical Hybrid Environment
for Reliability and Performance Modeling**

Summary of Research

03/01/1999 — 02/28/2003

Principal Investigator: Gianfranco Ciardo

Phone: +1-757-221-3478

Fax: +1-757-221-1717

e-mail: ciardo@cs.wm.edu

College of William and Mary, PO Box 8795, Williamsburg, VA 23187-8795

Distribution:

Grants Officer

NASA Langley Research Center, Mail Stop 126, Hampton, VA, 23681-2199

Plesent W. (Mike) Goode

NASA Langley Research Center, Mail Stop 152-D, Hampton, VA, 23681-2199

NASA Center for AeroSpace Information (CASI), Attn: Accessioning Department

7121 Standard Drive, Hanover, MD, 21076-1320

Mark Roberts

Office of Grants and Research Administration, College of William and Mary

Summary of Research for Grant Number: NAG-1-2168

An Advanced Hierarchical Hybrid Environment for Reliability and Performance Modeling

Principal Investigator: Gianfranco Ciardo

The key issue we intended to address in our proposed research project was the ability to model and study logical and probabilistic aspects of large computer systems. In particular, we wanted to focus mostly on automatic solution algorithms based on a state-space exploration as their first step, in addition to the more traditional discrete-event simulation approaches commonly employed in industry. One explicitly-stated goal was to

...extend by several orders of magnitude the size of models that can be solved exactly, using a combination of techniques:

- *Efficient exploration and storage of the state space using new data structures that require an amount of memory sublinear in the number of states.*
- *Exploitation of the existing symmetries in the matrices describing the system behavior using Kronecker operators.*

Not only we have been successful in achieving the above goals, but we exceeded them in many respects.

Accomplishments

Our research has substantially moved forward the boundaries of what models can be considered “solvable” by today’s standard.

For logical analysis, the use of decision diagrams for the storage of sets of states, Kronecker encodings for the storage of the transition relation, and the “saturation” algorithm for the generation of the state space has enormously increased the sizes of state spaces that can be analyzed. Depending on their structure and regularity, models with state spaces of 10^{20} — 10^{500} can be studied in a matter of minutes or at most hours.

For stochastic analysis, the use of Kronecker encodings and “matrix diagrams” has increased the size of solvable models by one to two orders of magnitude. Thus, it is now possible to compute exact stationary or transient performance measures on Markov models with up to 10^8 states.

Most importantly, perhaps, the work performed under this grant has opened up new exciting research directions in the field of modeling, while, at the same time, resulting in techniques that are already applicable today for the study of real systems. In addition to current and former students at William and Mary, the techniques we have introduced have been adopted and are in use by researchers at University of Illinois at Urbana Champaign, Carnegie Mellon University, University of California at Los Angeles, James Madison University, Supercomputer Center of the University of California at San Diego, AT&T Bell Labs, Polytechnic University Of Catalonia (Barcelona, Spain), University of Torino (Italy), Institute of Research on Communication and Cybernetics (Nantes, France), University of Reims (France), Technical University of Dresden (Germany), Brandenburg University of Technology at Cottbus (Germany), Technical University of Braunschweig (Germany), Technical University of Berlin (Germany), University of Magdeburg (Germany), and University of Edinburgh (UK).

The following section lists specific contributions supported by this grant, according to the publications where they appeared.

Publications acknowledging support by the grant

- [1] Discusses efficient data structures to encode very large Markov chains and provides fast algorithms to solve the corresponding models. Achieves at least one order of magnitude increase in the size of the models that can be stored and solved with respect to the best sparse-matrix techniques.
- [2] Surveys the state-of-the-art in advanced techniques for storing the reachability set and the transition rate matrix, with particular attention to the use of decision diagrams, Kronecker representations, and their interplay. Show that it is now possible to generate and store enormous reachability sets for logical analysis.
- [3] Proposes a family of new Kronecker-based techniques for the study of Markov models and compares them with existing ones, both traditional and Kronecker-based. The main contribution is the introduction of algorithms that use probability vectors of size equal to the “actual” state space instead of the much larger “potential” state space, providing great space and time savings.
- [4] Introduces a new approximation algorithm based on an exact representation of the state space, using decision diagrams, and of the transition rate matrix, using Kronecker algebra, for a Markov model decomposed into submodels. The algorithm provides exact results are obtained if the overall model has a product-form solution. Advantages of the method include good accuracy, low memory requirements, fast execution times, and a high degree of automation, since the only additional information required to apply it is a partition of the model into submodels. This is the first time an approximation algorithm has been proposed where knowledge of the exact state space is explicitly used.
- [5] Presents an algorithm to generate the state space of an asynchronous system using Multi-valued Decision Diagrams. In contrast to related work, the next-state function of a system is not encoded as a single Boolean function, but as cross-products of integer functions. This permits the application of various iteration strategies to build a system’s state space, including a new elegant strategy called “saturation”. In addition to performing several orders of magnitude faster than existing BDD-based state-space generators, the new algorithm requires peak memory that is often close to the final memory needed for storing the overall state space.
- [6] Considers two approaches to cope with the “state-space explosion”. Distributed algorithms that make use of the processors and memory overall available on a network of workstations can manage models with state spaces much larger than what is possible on a single workstation. A second approach, constituting a fundamental paradigm shift, is instead based on decision diagrams and related implicit data structures that efficiently encode the state space or the transition rate matrix of a model, provided that it has some structure to guide its decomposition; with these implicit methods, enormous sets can be managed efficiently.
- [7] **Keynote invited talk at PNPM’01.** Surveys techniques to cope with large state spaces, starting from early explicit methods, which require data structures of size proportional to the number of states or state-to-state transitions, then moving to implicit methods, which borrow ideas from symbolic model checking (binary decision diagrams) and numerical linear algebra (Kronecker operators) to drastically reduce the computational requirements. Introduces the “structural decomposition” approach. This method only requires to specify a partition of the places in the net and, combining decision diagrams and Kronecker operators with the new concepts of event locality and node saturation, achieves fundamental gains in both memory and time efficiency. At the same, the approach is applicable to a wide range of models.

- [8] Presents a novel stochastic Petri net formalism where both discrete and continuous phase-type firing delays can appear simultaneously in the same model. By capturing non-Markovian behavior in discrete or continuous time, as appropriate, the formalism affords higher modeling fidelity. Alone, discrete or continuous phase-type Petri nets have simple underlying Markov chains, but mixing the two complicates matters. In a mixed model where discrete-time transitions are synchronized, the underlying process is semi-regenerative and Markov renewal theory can be employed to formulate stationary or time-dependent solutions. Computational trade-offs between the so-called embedded and subordinate Markov chains can be employed to improve the overall solution efficiency.
- [9] Presents a new method for the symbolic construction of shortest paths in reachability graphs. The algorithm relies on a variant of edge-valued decision diagrams that supports efficient fixed-point iterations for the joint computation of both the reachable states and their distance from the initial states. Once the distance function is known, a shortest path from an initial state to a state satisfying a given condition can be easily obtained. Using a few representative examples, the algorithm is shown to be vastly superior, in terms of both memory and space, to alternative approaches that compute the same information, such as ordinary or algebraic decision diagrams. This algorithm has useful applications to counter-example generation in symbolic model checking.
- [10] Extends the “saturation” algorithm for symbolic state-space generation, which is characterized by the use of multi-valued decision diagrams, boolean Kronecker operators, event locality, and a special iteration strategy. Saturation outperforms traditional BDD-based techniques by several orders of magnitude in both space and time but, like them, assumes a priori knowledge of each submodel’s state space. Here, a new saturation algorithm is introduced that merges explicit local state-space discovery with symbolic global state-space generation. This relieves the modeler from worrying about the behavior of submodels in isolation, increasing the likelihood of a correct model and decreasing the specification burden on the user.
- [11] Thanks to a Kronecker encoding of the transition relation, event locality can be exploited and better fixed-point iteration strategies can be applied for state space generation. This paper extends these results to symbolic CTL model checking, resulting in orders-of-magnitude reductions for both execution times and memory consumption in comparison to well-established tools such as NuSMV.
- [12] Most of the research results described in this report have been implemented in the software tool SMART, which has been demonstrated at many conferences on verification, performance, and reliability evaluation. SMART has been widely distributed, is being used in the classroom, and is now over 100,000 lines of C++ code.

Supported students

The following students were supported at various times under this grant:

- Arun Mangalam (MS, Spring 2000)
- Michelle Griffith (MS, Fall 2000)
- Matthew Klinger (BA, Spring 2003)
- Robert Marmorstein (MS, Spring 2003)
- Radu Siminiceanu (PhD, expected Fall 2003)
- Raymond Plante (MS, expected Spring 2004)

References

- [1] Gianfranco Ciardo and Andrew S. Miner. A data structure for the efficient Kronecker solution of GSPNs. In Peter Buchholz, editor, *Proc. 8th Int. Workshop on Petri Nets and Performance Models (PNPM'99)*, pages 22–31, Zaragoza, Spain, September 1999. IEEE Comp. Soc. Press.
- [2] Gianfranco Ciardo and Andrew S. Miner. Structural approaches for SPN analysis. In Adrian Tenter, editor, *High Performance Computing 2000, Grand Challenges in Computer Simulation*, pages 345–356. SCS, April 2000.
- [3] Peter Buchholz, Gianfranco Ciardo, Susanna Donatelli, and Peter Kemper. Complexity of memory-efficient Kronecker operations with applications to the solution of Markov models. *INFORMS J. Comp.*, 12(3):203–222, 2000.
- [4] Andrew S. Miner, Gianfranco Ciardo, and Susanna Donatelli. Using the exact state space of a Markov model to compute approximate stationary measures. In Jim Kurose and Philippe Nain, editors, *Proc. ACM SIGMETRICS*, pages 207–216, Santa Clara, CA, June 2000. ACM Press.
- [5] Gianfranco Ciardo, Gerald Luetzgen, and Radu Siminiceanu. Saturation: An efficient iteration strategy for symbolic state space generation. In Tiziana Margaria and Wang Yi, editors, *Proc. Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, LNCS 2031, pages 328–342, Genova, Italy, April 2001. Springer-Verlag.
- [6] Gianfranco Ciardo. Distributed and structured analysis approaches to study large and complex systems. In Ed Brinksma, Holger Hermanns, and Joost-Pieter Katoen, editors, *Lectures on Formal Methods and Performance Analysis*, LNCS 2090, pages 344–374. Springer-Verlag, 2001.
- [7] Gianfranco Ciardo. What a structural world. In Reinhard German and Boudewijn Haverkort, editors, *Proc. 9th Int. Workshop on Petri Nets and Performance Models (PNPM'01)*, pages 3–16, Aachen, Germany, September 2001. IEEE Comp. Soc. Press. Invited keynote paper.
- [8] Robert L. Jones and Gianfranco Ciardo. On phased delay stochastic Petri nets: Definition and an application. In Reinhard German and Boudewijn Haverkort, editors, *Proc. 9th Int. Workshop on Petri Nets and Performance Models (PNPM'01)*, pages 165–174, Aachen, Germany, September 2001. IEEE Comp. Soc. Press.
- [9] Gianfranco Ciardo and Radu Siminiceanu. Using edge-valued decision diagrams for symbolic generation of shortest paths. In Mark D. Aagaard and John W. O'Leary, editors, *Proc. Fourth International Conference on Formal Methods in Computer-Aided Design (FMCAD)*, LNCS 2517, pages 256–273, Portland, OR, USA, November 2002. Springer-Verlag.
- [10] Gianfranco Ciardo, Robert Marmorstein, and Radu Siminiceanu. Saturation unbound. In Hubert Garavel and John Hatcliff, editors, *Proc. Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, LNCS 2619, pages 379–393, Warsaw, Poland, April 2003. Springer-Verlag.
- [11] Gianfranco Ciardo and Radu Siminiceanu. Structural symbolic CTL model checking. In *Computer Aided Verification (CAV'03)*, LNCS, Boulder, CO, USA, July 2003. Springer-Verlag. To appear.
- [12] Gianfranco Ciardo, Robert L. Jones, Robert M. Marmorstein, Andrew S. Miner, and Radu Siminiceanu. SMART: Stochastic Model-checking Analyzer for Reliability and Timing. In *Proc. Int. Conf. on Dependable Systems & Networks (DSN)*, page 545, Washington, D.C., USA, June 2002. IEEE Press.